

1 WE CLAIM:

1 A network switch for resolving requests from a plurality of host initiators by scheduling
2 access to a plurality of disk storage devices, the network switch comprising:
3 (a) a switched fabric comprising a plurality of switching elements, each switching
4 element comprising:
5 a plurality of bi-directional switched fabric ports; and
6 a control input connected to receive switch control data for selectively configuring
7 the switching element in order to interconnect the bi-directional switched
8 fabric ports;
9 (b) a memory for storing a routing and scheduling program; and
10 (c) a microprocessor, responsive to the requests, for executing the steps of the routing
11 and scheduling program to generate the switch control data to transmit scheduled
12 requests through the bi-directional switched fabric ports,
13 wherein:
14 at least one of the plurality of switching elements comprises a disk storage interface
15 for connecting to a selected one of the disk storage devices;
16 the microprocessor for scheduling access to the plurality of disk storage devices
17 through the disk storage interface;
18 the disk storage interface for receiving scheduling data from the selected one of the
19 storage devices;
20 the memory for receiving the scheduling data via the bi-directional switched fabric
21 ports of a selected number of the switching elements; and
22 the scheduling data is processed according to a priority such that the selected
23 switching elements transfer the scheduling data through the bi-directional
24 switched fabric ports before transferring data associated with the scheduled
25 requests.

- 1 2. The network switch as recited in claim 1, wherein the at least one switching element
2 further comprises a disk storage device connected to the disk storage interface.
- 1 3. The network switch as recited in claim 1, wherein:
2 (a) each disk storage device comprises a disk and a head; and
3 (b) the scheduling data comprises a radial location of the head relative to the disk within
4 each disk storage device.

B1
1 4. The network switch as recited in claim 3, wherein the scheduling data further comprises a
2 circumferential location of the head relative to the disk within each disk drive.
- 1 5. The network switch as recited in claim 1, wherein the switching elements further
2 comprise a plurality of virtual lanes, wherein:
3 (a) at least one of the virtual lanes is reserved for transferring data associated with the
4 scheduled requests;
5 (b) at least one of the virtual lanes is reserved for transferring the scheduling data; and
6 (c) the virtual lane for transferring the scheduling data comprises a higher priority than
7 the virtual lane for transferring the data associated with the scheduled requests.
- 1 6. The network switch as recited in claim 1, wherein the scheduling data is communicated to
2 the memory through the bi-directional switched fabric ports according to an isochronous
3 protocol.

1 7. A method of resolving requests from a plurality of host initiators by scheduling access to
2 a plurality of disk storage devices connected to a network switch, the network switch
3 comprising a switched fabric comprising a plurality of switching elements, the method
4 comprising the steps of:

- 5 (a) transmitting through the switching elements scheduling data from the plurality of
6 disk storage devices to a memory;
7 (b) evaluating the scheduling data in order to schedule the requests from the host
8 initiators; and
9 (c) transmitting data associated with the scheduled requests through the switching
10 elements to the plurality of disk storage devices,

11 wherein the scheduling data is processed according to a priority such that the
12 switching elements transfer the scheduling data before transferring data associated
13 with the scheduled requests.

1 8. The method as recited in claim 7, wherein:

- 2 (a) each disk storage device comprises a disk and a head; and
3 (b) the scheduling data comprises a radial location of the head relative to the disk within
4 each disk storage device.

1 9. The method as recited in claim 8, wherein the scheduling data further comprises a
2 circumferential location of the head relative to the disk within each disk drive.

1 10. The method as recited in claim 7, wherein the switching elements further comprise a
2 plurality of virtual lanes, wherein:

- 3 (a) at least one of the virtual lanes is reserved for transferring data associated with the
4 scheduled requests;
5 (b) at least one of the virtual lanes is reserved for transferring the scheduling data; and

6 (c) the virtual lane for transferring the scheduling data comprises a higher priority than
7 the virtual lane for transferring the data associated with the scheduled requests.

1 11. The method as recited in claim 7, wherein the scheduling data is communicated to the
2 memory through the switching elements according to an isochronous protocol.